



05/17/05

ZW
AFF

Express Mail No. EV544475406US

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of:
MOULTON

Serial No.: 09/782,532

Filed: February 13, 2001

For: **SYSTEM AND METHOD FOR
DISTRIBUTED MANAGEMENT
OF DATA STORAGE**

Confirmation No. 7507

Art Unit: 2151

Examiner: PHILLIPS, HASSAN

Customer No. 23235

Docket No. UND011

CERTIFICATE OF MAILING BY EXPRESS MAIL

MAIL STOP APPEAL BRIEF – PATENTS
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

The undersigned hereby certifies that the following documents:

- Appellant's Brief Under 37 CFR 1.192 (6 pages + 8 pages of claims Appendix A3 copies);
- Check in the amount of \$500.00;
- Certificate of Mailing by Express Mail; and
- Return Receipt Postcard

relating to the above application, were deposited as "Express Mail", Mailing Label No. EV544475406US with the U.S. Postal Service, addressed to Mail Stop Appeal Brief – Patents, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

5/16/04
Date

5/16/04
Date

Stuart Langley
Mailer

Stuart Langley
Stuart T. Langley, Reg. No. 33,940
HOGAN & HARTSON LLP
One Tabor Center
1200 17th Street, Suite 1500
Denver, Colorado 80202
(720) 406-5335 Tel
(303) 899-7333 Fax



Attorney Docket No. UND011
Client/Matter No. 83208.0014
Express Mail No. VIA FACSIMILE

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES**

In re Application of:

MOULTON
Serial No.: 09/782,532

Filed: February 13, 2001

For: **SYSTEM AND METHOD FOR
DISTRIBUTED MANAGEMENT
OF DATA STORAGE**

Confirmation No. 7507

Art Unit: 2151

Examiner: PHILLIPS, HASSAN

Customer No. 25235

Docket No. UND011

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPELLANT'S BRIEF UNDER 37 CFR 1.192

I. Real Party in Interest

Avamar Technologies, Inc.
135 Technology Drive, Suite 100
Irvine, CA 92618 USA

II. Related Appeals and Interferences

No other appeals or interferences are currently known to Appellant that will directly affect, be directly affected by, or have a bearing on the decision to be rendered by the Board of Patent Appeals and Interferences in the present appeal.

05/18/2005 MAHMED1 00000016 09782532

01 FC:1402

500.00 OP

III. Status of Claims

Claims 1, 3-10, 12-21, 23, 24, 26-28, 33-37, and 39-54 remain in the application. No claims have been allowed. However, as set out below, one of the issues is whether claims 1, 3-10, 12-21, 23, 24, 26-28, 33-37 and 44-51 are actually rejected.

Claims 52 and 53 were rejected under 35 U.S.C. 102 based upon Carter.

Claims 42-43 were rejected under 35 U.S.C. 103 based upon Carter in view of Gershman.

IV. Status of Amendments

The claim amendments filed on January 12, 2005 have been entered. Claims 1, 3-10, 12-21, 23, 24, 26-28, 33-37, and 39-54 as currently amended are provided in the Appendix of Claims on Appeal.

V. Summary of the Invention

The present invention involves a data storage system such as shown in Fig. 1 and described at pages 9-10 of the specification. The system implements storage management functionality in a distributed manner using storage management mechanisms such as 106 shown in Fig. 1 and initially described in the paragraph bridging pages 12 and 13. Preferably, the storage management system comprises a plurality of instances of storage management processes 106 where the instances are physically distributed such that failure or unavailability of any given instance or set of instances will not impact the availability of stored data as set out at page 14, lines 5-14.. The storage management process instances communicate with each other to store data in a distributed, collaborative fashion with no centralized control of the system (see, e.g., page 13, line 30).

In a particular implementation, the present invention involves systems and methods for distributing data with parity (e.g., redundancy) over a large geographic and topological area in a network architecture. Data is transported to, from, and between nodes using network connections 101 rather than bus connections as

described, for example, at page 12, lines 1-28 of the specification. The network data distribution relaxes or removes limitations on the number of storage devices and the maximum physical separation between storage devices that limited prior fault-tolerant data storage systems and methods. The present invention allows data storage to be distributed over larger areas (e.g., the entire world in the case of the system shown in Fig. 1), thereby mitigating outages from localized problems such as network failures, power failures, as well as natural and man-made disasters.

VI. Issues

- A. **On what grounds are claims 1, 3-10, 12-21, 23, 24, 26-28, 33-37 and 44-51 rejected?**
- B. **Is the Final Rejection Proper?**
- C. **Is the rejection of claims 52 and 53 Proper?**
- D. **Is the rejection of claims 42-43 Proper?**

VII. Grouping of Claims

- A. **With respect to issue A, claims 1, 3-10, 12-21, 23, 24, 26-28, 33-37 and 44-51 stand together.**
- B. **With respect to issue B, all of the claims stand together.**
- C. **With respect to issue C, claims 52 and 53 stand together.**
- D. **With respect to issue C, claims 42-43 stand together.**

VIII. Argument

- A. **Are claims 1, 3-10, 12-21, 23, 24, 26-28, 33-37 and 44-51 rejected?**

The final Office action does not state any ground of rejection for claims 1, 3-10, 12-21, 23, 24, 26-28, 33-37 and 44-51. Although it appears that the Office action responds in negatively to the arguments made in the response filed on August 3, 2004, there is only a statement that “the references supplied by the examiner in the previous Office action covers claims 1, 3-10, 12-21, 23, 24, 26-28, 33-37 and 44-51. Unfortunately, “covers” is not a clear statutory ground for rejecting claims and applicant is left uncertain as to whether a rejection under 35 U.S.C. 102, 35 U.S.C. 103, or some other statute is intended.

The Advisory Action does not make any attempt to clarify the status of the claims. Because the Office Action leaves the status of claims 1, 3-10, 12-21, 23, 24, 26-28, 33-37 and 44-51 ambiguous, finality of the Office Action was (and is) believed to be improper and should be overturned by the Board.

Although a rejection is not clearly stated, it is believed that Carter does not show or suggest the use of processes that implement RAID-type distribution as called for in claim 1 as Carter does not hint at the nature of any processes used to implement the RAID disk mentioned in col. 16. Further, Carter does not show or suggest a parity scheme. The brief mention of RAID does not suggest parity as several RAID configurations do not use parity. Likewise, the replication or copying of data in multiple locations is a different concept than parity protection. With respect to communicating operational state information, it is maintained that even if Carter teaches communicating state information, Carter is limited to information describing the state of the data, not the state of the device(s) in which the data exists. These are very different types of state information.

The final Office action asserts that “fault tolerance” in Carter is the same as determining when a fault is likely. However, Carter is clear throughout that fault tolerance refers to behavior that occurs after a fault, not before a fault. In actuality, if one could anticipate a fault and migrate data prior to a fault as called for in some of Applicant’s claims, “fault tolerance” as that term is typically used would no longer be necessary (assuming the fault anticipation operated perfectly, of course).

Accordingly, it is respectfully requested that the Board clarify whether claims 1, 3-10, 12-21, 23, 24, 26-28, 33-37 and 44-51 are rejected, and, if they are rejected to reverse that rejection.

B. Is the Final Rejection Proper?

Because the Office Action leaves the status of claims 1, 3-10, 12-21, 23, 24, 26-28, 33-37 and 44-51 ambiguous, finality of the Office Action is believed to be improper and should be overturned by the Board. Further, the Advisory Action justifies the finality of the office action because of amendments made to claim 42. However, claim 42 was not amended. It is not entirely clear whether the response

filed on January 12, 2005 was even considered. Should the Board decision require further prosecution, it is respectfully requested that the finality of the Office action be overturned.

C. Rejection of claims 52 and 53 under 35 U.S.C. §103 is Improper.

Claims 52 and 53 are distinct with respect to Carter for at least the same reasons as claim 1. Moreover, Carter does not show or suggest the RAID-type distribution comprises managing redundancy operations across the plurality of network-accessible devices. To the extent Carter teaches anything about RAID, it refers to a single disk or single device that provides RAID storage and does not show or suggest implementing such processes across several disks or devices. With respect to claim 53, the cited portion of Carter does not show or suggest anything about data striping, data mirroring, parity data distribution, parity data mirroring, and data entry as N-separated secrets.

Based on the above remarks, Appellant respectfully requests that the rejection of claims 52-53 be reversed.

D. Rejection of claim 42-43 under 35 U.S.C. §103 is Improper.

Claims 42 and 43 are believed to be distinct over Carter for the reasons stated in the August 3, 2004 response. Specifically, Independent claim 42 calls for, among other things, providing a plurality of network accessible storage devices and “monitoring the data storage for faults” with storage management processes by having at least a portion of the plurality of network accessible storage devices transmitting heartbeat messages. Carter does not teach monitoring the devices. While Gershman is relied on to show a heartbeat message, Gershman does not supply the basic deficiency of Carter in that Gershman does not teach monitoring a plurality of network accessible storage devices.

Based on the above remarks, Appellant requests that the rejection of claims 42 and 43 be reversed.

E. Rejection of claim 54 under 35 U.S.C. §103 is Improper.

Claim 54 depends from claim 21 and is believed to be distinct with respect to Carter for the same reasons as claim 21. Thompson does not supply the deficiencies of Carter. Specifically, Thompson does not show or suggest a method of data storage management that calls for “communicating state information between the plurality of storage devices.” For at least these reasons claim 54 is believed to be allowable over the relied on references either alone or in combination.

Based on the above remarks, Appellant requests that the rejection of claim 54 be reversed.

Conclusion

In view of all of the above, claims 1, 3-10, 12-21, 23, 24, 26-28, 33-37, and 39-54 are believed to be allowable and the case in condition for allowance. Appellant respectfully requests that the Examiner’s rejections be reversed for all pending claims.

Respectfully submitted,

Date: May 16, 2005



Stuart T. Langley, Reg. No. 33,940
HOGAN & HARTSON LLP
One Tabor Center
1200 17th Street, Suite 1500
Denver, Colorado 80202
Phone: (720) 406-5335
Fax: (720) 406-5301



IX. APPENDIX OF CLAIMS ON APPEAL

Claim 1: A data storage management system comprising:
at least one network-accessible storage device capable of storing data, wherein
the at least one network-accessible device capable of storing data comprises a
plurality of network-accessible devices capable of storing data, some of which are
located at distinct network nodes;
a plurality of network-accessible devices configured to implement storage
management processes; and
a communication system enabling the storage management processes to
communicate with each other;
wherein the storage management processes comprise processes for storing
data to the at least one network-accessible device; and
wherein the processes for storing data comprise processes that implement a
RAID-type distribution across the plurality of network-accessible devices.

Claim 2 (canceled)

Claim 3: The data storage system of claim 1 wherein the storage management
processes comprise processes for serving data from the at least one network accessible
storage device.

Claim 4: The data storage system of claim 1 wherein the at least one storage
device comprises a RAID storage system.

Claim 5: The data storage system of claim 1 wherein the at least one storage
device comprises a computer with direct attached storage (DAS) selected from the
group consisting of magnetic hard disk, magneto-optical, optical disk, digital optical
tape, holographic storage, quantum storage, and atomic force probe storage.

Claim 6: The data storage system of claim 1 wherein the plurality of storage
devices comprises a peer-to-peer network of storage devices, each storage device
having means for communicating state information with other storage devices, at least

one storage device comprising means for receiving storage requests from external entities, and at least one storage device comprising means for causing read and write operations to be performed on others of the storage devices.

Claim 7: The data storage system of claim 1 wherein the communication system comprises a TCP/IP over Ethernet network.

Claim 8: The data storage system of claim 1 wherein the communication system comprises Gigabit Ethernet network.

Claim 9: The data storage system of claim 1 wherein the communication system comprises a Fibre Channel fabric.

Claim 10: The data storage system of claim 1 wherein the communication system comprises a wireless network.

Claim 11 (canceled)

Claim 12: The data storage system of claim 1 wherein the processes for storing data comprise processes that implement an n-dimensional parity scheme for data elements across the plurality of network accessible devices.

Claim 13: The data storage system of claim 12 wherein the processes for storing parity data expand or contract a size of a parity group associated with each of the data elements associated with the n-dimensional parity scheme to whatever extent is desired.

Claim 14: The data storage system of claim 12 wherein the storage management processes further comprise processes for recovery of data when one or more of the network-accessible storage devices is unavailable.

Claim 15: The data storage system of claim 12 wherein the storage management processes further comprise processes for access to stored data when one or more of the network accessible storage devices are not desirable data sources for

reasons including but not limited to efficiency, performance, network congestion, and security.

Claim 16: The data storage system of claim 1 wherein the plurality of network-accessible devices configured to implement storage management processes further comprise commercial off-the-shelf computer systems implementing a common operating system.

Claim 17: The data storage system of claim 1 wherein the plurality of network-accessible devices configured to implement storage management processes further comprise commercial off-the-shelf computer systems implementing a heterogeneous set of operating systems.

Claim 18: The data storage system of claim 1 wherein the storage management processes comprise processes for implementing greater than two dimensions of parity.

Claim 19: The data storage system of claim 1 wherein the processes for storing data comprise processes that store parity and/or mirror data across more than one of the plurality of network-accessible storage devices.

Claim 20: The data storage system of claim 1 wherein the storage management processes comprise processes for adding and removing additional storage capacity to individual storage devices and the system as a whole.

Claim 21: A method of data storage management comprising the acts of:
providing at least one network-accessible storage device capable of storing data, wherein the at least one network-accessible device capable of storing data comprises a plurality of network-accessible storage devices capable of storing data, some of which are located at distinct network nodes;
implementing a plurality of storage management process instances;
communicating storage messages between the storage management process instances;

storing data to the at least one network-accessible device under control of at least one instance of the storage management processes; and

implementing a peer-to-peer network between the plurality of storage devices;

communicating state information for the plurality of network-accessible storage devices between the plurality of network-accessible storage devices; and

performing read and write operations using the plurality of storage devices.

Claim 22 (canceled)

Claim 23: The method of claim 21 further comprising serving data from the at least one network accessible storage device.

Claim 24: The method of claim 21 wherein the step of storing data to the at least one storage device comprises storing the data in a RAID-like fashion.

Claim 25 (canceled)

Claim 26: The method of claim 21 wherein the step of storing data comprises storing data using a RAID-type distribution across the plurality of network-accessible storage devices.

Claim 27: The method of claim 21 wherein the act of storing data comprises storing parity and/or mirror data across more than one of the plurality of network-accessible storage devices.

Claim 28: The method of claim 21 wherein the storage management process instances further comprise processes for recovery of data when one or more of the network-accessible storage devices is unavailable.

Claims 29-32 (canceled)

Claim 33: A method of data storage management comprising the acts of:
providing a plurality of network-accessible storage devices each capable of storing data;
implementing a plurality of storage management process instances;

communicating storage messages between the storage management process instances; and

identifying two or more storage devices associated with the data to be stored;
determining parity data for the data to be stored;
storing the data and/or the parity data across the two or more storage devices;
retrieving the stored data;
verifying the correctness of the stored data using the parity data; and
upon detection of an error in the retrieved data, retrieving a correct version of the data using the parity data.

Claim 34: The method of claim 33 wherein the parity data comprises an error checking and correcting code.

Claim 35: The method of claim 33 wherein the parity data comprises a mirror copy of the data to be stored.

Claim 36: The method of claim 33 wherein the parity data is stored in a single network storage node and the unit of data is stored in two or more network storage nodes.

Claim 37: The method of claim 33 wherein the parity data is distributed across multiple storage nodes.

Claim 38 (canceled)

Claim 39: The method of claim 33 further comprising:
attempting to retrieve the stored data;
detecting unavailability of one of the two or more network storage nodes; and
in response to detecting unavailability, reconstructing a correct version of the data using the parity data.

Claim 40: The system of claim 33 wherein the act of storing the data comprises distributing non-identical but logically equivalent data in a storage node.

Claim 41: The system of claim 33 further comprising storing lossy equivalent data in a storage node.

Claim 42: A method of data storage management comprising the acts of : providing a plurality of network accessible storage devices capable of storing data;

implementing a plurality of storage management process instances; communicating storage messages between the plurality of storage management processes;

storing data to the plurality of network accessible storage devices under control of the plurality of storage management processes;

adding and subtracting data storage capacity to and from the data storage under control of the plurality of storage management processes without affecting accessibility of the data storage; and

monitoring the data storage for faults by means of the plurality of storage management processes, wherein the monitoring comprises at least a portion of the plurality of network accessible storage devices transmitting heartbeat messages.

Claim 43: The method of claim 42 further comprising:
compensating for the faults by manipulating the data storage under control of the plurality of storage management processes without affecting accessibility of the data storage.

Claim 44: A method of data storage management comprising the acts of: providing a plurality of network-accessible storage devices each capable of storing data;

implementing a plurality of storage management process instances; and
communicating storage messages between the storage management process instances, wherein any of the storage management process instances is capable of storage allocation and deallocation across the plurality of network-accessible storage devices;

wherein the storage management processes are configured to migrate data amongst the storage devices using the storage messages preemptively when a fault condition in at least one of the storage devices is determined to be likely.

Claim 45: The method of claim 44 wherein the storage allocation management processes are configured to use the storage messages to reconstruct data stored in a failed one of the storage devices.

Claim 46: The method of claim 44 wherein the storage management processes are configured to migrate data amongst the storage devices using the storage messages in response to a detected fault condition in at least one of the storage devices.

Claim 47 (canceled)

Claim 48: The method of claim 44 wherein the plurality of storage devices comprises an arbitrarily large number of storage devices.

Claim 49: The method of claim 44 further comprising:
associating parity information with a data set;
storing the parity information in at least some of the storage devices; and
serving data requests corresponding to the data set by accessing the parity information associated with the data set.

Claim 50: The method of claim 44 further comprising:
storing a data set in a plurality of the data storage devices using the storage management processes;
serving data requests corresponding to the data set by accessing the plurality of data storage devices in parallel.

Claim 51: The method of claim 44 further comprising encrypting storage messages before communicating.

Claim 52: The method of claim 1, wherein the RAID-type distribution comprises managing redundancy operations across the plurality of network-accessible devices.

Claim 53: The method of claim 1, wherein the RAID-type distribution comprises one or more functionalities selected from the group consisting of data striping, data mirroring, parity data distribution, parity data mirroring, and data entry as N-separated secrets.

Claim 54: The method of claim 21, wherein the state information comprises access speed, transfer rate, network locality, physical locality, interconnectedness, security, reliability, political domain, capacity, or cost.